

## CLAIMS

1. A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile  
5 storage media, the method comprising:

maintaining a record predictive of locations to which data are to be written on the primary storage subsystem by a host processor;

receiving the data from the host processor at the  
10 primary storage subsystem to be written to a specified location on the first non-volatile storage media;

if the specified location is not included in the record, updating the record responsively to the specified location;

15 signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, if the specified location was not included in the record, responsively to updating the record;

20 copying the data from the primary storage subsystem to the secondary storage subsystem; and

storing the data in the specified location on both the first and second non-volatile storage media.

25 2. The method according to claim 1, wherein copying the data comprises transmitting the data between mutually-remote sites over a communication link between the sites.

3. The method according to claim 1, wherein copying the  
30 data comprises creating a mirror on the secondary storage

subsystem of the data received by the primary storage subsystem.

4. The method according to claim 3, and comprising,  
5 upon occurrence of a failure in the primary storage subsystem, configuring the secondary storage subsystem to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

10

5. The method according to claim 3, and comprising,  
upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record,  
a portion of the data from the secondary storage  
15 subsystem to the primary storage subsystem for storage on the primary storage subsystem.

6. The method according to claim 1, wherein maintaining and updating the record comprise marking respective bits  
20 in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

7. The method according to claim 1, wherein maintaining the record comprises storing the record on the first non-volatile storage media, and wherein updating the record  
25 comprises modifying the record that is stored on the first non-volatile storage media.

8. The method according to claim 7, wherein modifying the record comprises:

5     comparing the specified location to a copy of the record held in a volatile memory on the primary storage subsystem;

      modifying the copy of the record so that at least the specified location is included in the copy of the record; and

10     destaging the modified copy of the record to the first non-volatile storage media.

9. The method according to claim 8, wherein the record is not modified on the first non-volatile storage media responsively to receiving the data as long as the  
15     specified location to which the data are to be written is included in the record.

10. The method according to claim 7, wherein modifying the record comprises adding a plurality of locations,  
20     including the specified location, to the record.

11. The method according to claim 1, wherein updating the record comprises predicting one or more further locations to which the host processor is expected to  
25     write the data in a subsequent write operation, and adding the one or more further locations to the record.

12. The method according to claim 11, wherein predicting the one or more further locations comprises selecting a

predetermined number of consecutive locations in proximity to the specified location.

13. The method according to claim 11, wherein  
5 maintaining the record comprises recording the locations to which the data are written using an object-based storage technique, and wherein predicting the one or more further locations comprises choosing the one or more further locations based on a logical connection between  
10 storage objects.

14. The method according to claim 1, wherein updating the record comprises removing one or more locations, other than the specified location, from the record, so as  
15 to limit a size of the record.

15. The method according to claim 14, wherein removing the one or more locations comprises receiving an acknowledgment from the secondary storage subsystem that  
20 the data have been stored in the one or more locations on the second non-volatile storage media, and removing the one or more locations from the record responsively to the acknowledgment.

25 16. The method according to claim 14, wherein removing the one or more locations comprises identifying the locations at which the first and second non-volatile storage media contain substantially identical data, and

selecting for removal one of the identified locations that was least-recently added to the record.

17. A data storage system, comprising:

5 a primary storage subsystem, which comprises first non-volatile storage media; and

a secondary storage subsystem, which comprises second non-volatile storage media,

wherein the primary storage subsystem is arranged to  
10 receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, which is arranged to store the data in the  
15 specified location on the second non-volatile storage media, and

wherein the primary storage subsystem is arranged to maintain a record predictive of locations to which data are to be written on the primary storage subsystem by the  
20 host processor, and upon receiving the data from the host processor, to update the record responsively to the specified location if the specified location is not included in the record, and to signal the host processor that the data have been stored in the data storage system  
25 responsively to receiving the data and, if the specified location was not included in the record, responsively to updating the record.

18. The system according to claim 17, wherein the first  
30 and second non-volatile storage media are located at mutually-remote sites, and wherein at least one of the

primary and secondary storage subsystems is arranged to transmit the data over a communication link between the sites.

5 19. The system according to claim 17, wherein the secondary storage subsystem is arranged to mirror the data held by the primary storage subsystem.

10 20. The system according to claim 19, wherein upon occurrence of a failure in the primary storage subsystem, the secondary storage subsystem is configurable to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

15

21. The system according to claim 19, wherein upon recovery of the system from a failure of the primary storage subsystem, the secondary storage subsystem is arranged to convey, responsively to the record, a portion  
20 of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

22. The system according to claim 17, wherein the record  
25 comprises a bitmap, and wherein the primary storage subsystem marks respective bits in the bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

23. The system according to claim 17, wherein the primary storage subsystem is arranged to store and update the record on the first non-volatile storage media.

5

24. The system according to claim 23, wherein the primary storage subsystem comprises a volatile memory and is arranged to hold a copy of the record in the volatile memory, and to update the record by modifying the copy of  
10 the record, and destaging the modified copy of the record to the first non-volatile storage media.

25. The system according to claim 24, wherein the record is not modified on the first non-volatile storage media  
15 responsively to receiving the data as long as the specified location to which the data are to be written is included in the record.

26. The system according to claim 23, wherein the  
20 primary storage subsystem is arranged, when the specified location is not included in the record, to update the record in the first non-volatile storage media by adding a plurality of locations, including the specified location, to the record.

25

27. The system according to claim 17, wherein the primary storage subsystem is arranged, when the specified location is not included in the record, to predict one or more further locations to which the host processor is

expected to write the data in a subsequent write operation, and to add both the specified location and the one or more further locations to the record.

5 28. The system according to claim 27, wherein the one or more further locations predicted by the primary storage subsystem comprise a predetermined number of consecutive locations in proximity to the specified location.

10 29. The system according to claim 27, wherein the primary storage subsystem is arranged to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

15 30. The system according to claim 17, wherein the primary storage subsystem is arranged, upon updating the record, to remove one or more locations, other than the specified location, from the record, so as to limit a size of the record.

20

31. The system according to claim 30, wherein the secondary storage subsystem is arranged to transmit an acknowledgment to the primary storage subsystem indicating that the data have been stored in the one or  
25 more locations on the second non-volatile storage media, and wherein the primary storage subsystem is arranged to remove the one or more locations from the record responsively to the acknowledgment.



32. The system according to claim 30, wherein the primary storage subsystem is arranged to identify the locations at which the first and second non-volatile storage media contain substantially identical data, and  
5 to remove from the record one of the identified locations that was least-recently added to the record.

33. A computer software product for use in a data storage system including primary and secondary storage  
10 subsystems, which include respective first and second control units and respective first and second non-volatile storage media, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first  
15 and second control units, cause the first control unit to receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage  
20 subsystem, and cause the second control unit to store the data in the specified location on the second non-volatile storage media,

wherein the instructions further cause the first control unit to maintain a record predictive of locations  
25 to which data are to be written on the primary storage subsystem by the host processor, and upon receiving the data from the host processor, to update the record responsively to the specified location if the specified location is not included in the record, and to signal the  
30 host processor that the data have been stored in the data storage system responsively to receiving the data and, if

the specified location was not included in the record,  
responsively to updating the record.

34. The product according to claim 33, wherein the first  
5 and second non-volatile storage media are located at  
mutually-remote sites, and wherein the instructions cause  
at least one of the first and second control units to  
transmit the data over a communication link between the  
sites.

10

35. The product according to claim 33, wherein the  
instructions cause the first and second control units to  
mirror the data held by the primary storage subsystem on  
the secondary storage subsystem.

15 36. The product according to claim 35, wherein the  
instructions cause the secondary storage subsystem, upon  
occurrence of a failure in the primary storage subsystem,  
to serve as the primary storage subsystem so as to  
receive further data from the host processor to be stored  
20 by the data storage system.

37. The product according to claim 35, wherein upon  
recovery of the system from a failure of the primary  
storage subsystem, the instructions cause the second  
25 control unit to convey, responsively to the record, a  
portion of the data from the second non-volatile storage  
media to the primary storage subsystem for storage on the  
first non-volatile storage media.

38. The product according to claim 33, wherein the record comprises a bitmap, and wherein the instructions cause the first control unit to mark respective bits in the bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

39. The product according to claim 33, wherein the instructions cause the first control unit to store and update the record on the first non-volatile storage media.

40. The product according to claim 39, wherein the instructions cause the first control unit to hold a copy of the record in a volatile memory of the primary storage subsystem, and to update the record by modifying the copy of the record, and destaging the modified copy of the record to the first non-volatile storage media.

41. The product according to claim 40, wherein the instructions cause the first control unit not to modify the record on the first non-volatile storage media responsively to receiving the data as long as the specified location to which the data are to be written is included in the record.

42. The product according to claim 39, wherein the instructions cause the first control unit, when the specified location is not included in the record, to

update the record in the first non-volatile storage media by adding a plurality of locations, including the specified location, to the record.

5 43. The product according to claim 33, wherein the instructions cause the first control unit, when the specified location is not included in the record, to predict one or more further locations to which the host processor is expected to write the data in a subsequent  
10 write operation, and to add both the specified location and the one or more further locations to the record.

44. The product according to claim 43, wherein the one or more further locations predicted by the first control  
15 unit comprise a predetermined number of consecutive locations in proximity to the specified location.

45. The product according to claim 43, wherein the instructions cause the first control unit to maintain the  
20 record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

46. The product according to claim 33, wherein the instructions cause the first control unit, upon updating  
25 the record, to remove one or more locations, other than the specified location, from the record, so as to limit a size of the record.

47. The product according to claim 46, wherein the instructions cause the second control unit to transmit an acknowledgment to the primary storage subsystem indicating that the data have been stored in the one or  
5 more locations on the second non-volatile storage media, and further cause the first control unit to remove the one or more locations from the record responsively to the acknowledgment.

10 48. The product according to claim 46, wherein the instructions cause the first control unit to identify the locations at which the first and second non-volatile storage media contain substantially identical data, and to remove from the record one of the identified locations  
15 that was least-recently added to the record.